
Rule WLM151: Server service class delays (multiple transaction service classes)

Finding: CPExpert has identified delays for the server service class that provided service to a subsystem transaction service class. The server provided service to more than one subsystem transaction service class, and CPExpert prorates the service provided to the different transaction service classes.

Impact: This finding is provided for information purposes.

Logic flow: The following rules cause this rule to be invoked:

- Rule WLM104: Subsystem Service Class did not achieve average response goal
- Rule WLM105: Subsystem Service Class did not achieve percentile response goal

Discussion: When CPExpert produces Rule WLM104 or Rule WLM105 to indicate that a subsystem service class did not achieve its performance goal, the logic of these rules tries to identify the cause of the delay. The cause of the delay initially is analyzed from the "served" service class view. Please refer to Rule WLM120 to Rule WLM132 for a discussion of the delays from the served service class.

After analyzing the **served** service class delays, CPExpert identifies the **server** service class. The server service class normally will be one or more CICS regions or IMS regions.

CPExpert analyzes the following possible delays to response time¹:

- **CPU Using delay**
- **Denied CPU delay**
- **CPU Capping delay**
- **Swap-in delay**
- **MPL delay**

¹Please see Section 4 (Chapter 3.3) for a description of these delays.

-
- **Page-in delay**
 - **I/O delay**
 - **Unknown delay**

If the server service class provides service to more than one transaction service class, CPExpert must apportion the resource utilization and delays to the different transaction service classes. The apportioning is done based on the value of SMF variable R723SCS# for each transaction service class.

CPExpert produces Rule WLM151 to provide a summary of the delay for the server service class.

The output from Rule WLM151 does not contain the MPL delay or swap-in delay. In most environments, server service classes are non-swappable and the MPL delay and swap-in delay columns would always show zero. Consequently, CPExpert does not clutter up the output with columns that almost always would be zero. However, CPExpert **does** analyze these delays if any are non-zero.

CPExpert produces Rule WLM151 to show the delays to the server service class. More than one server service class might serve the subsystem transaction service class that missed its performance goal. In this case, CPExpert produces multiple Rule WLM151 findings - one for each server service class. CPExpert then analyzes the delays for each server service class.

Additionally, a server service class might serve more than one transaction subsystem service class (in fact, this is the more common case). For example, a CICS region often will serve several transaction service classes composed of CICS transactions. In this case, CPExpert must apportion the resources used and delays encountered among the transaction service classes being served. The resources and delays encountered are reported by Rule WLM151.

Rule WLM151 also shows the percent of service provided the transaction service class missing its performance goal, relative to the service provided by the server to all transaction service classes.

- There is no information in SMF Type 72 records that shows how much of the response time of the **served** service class (e.g., the IMS Service Class) could be attributed to delays in the **individual servers** (e.g., CICS RGN).

- If the individual servers serve more than one service class, there is information in the SMF Type 72 records to show how many **times** an address space in the server was observed to be providing service to the served service class.

The WLM counts each time the server issues the IWMRPT macro to indicate that a transaction has completed. This count lets the WLM know how many times the server (e.g., a CICS region) provided service to the served service class (e.g., CICS transactions).

Additionally, every 250 milliseconds, the WLM samples server service classes to see which served service classes they are serving. The sampling process ensures that the WLM keeps track of service provided to long-running transactions.

SMF field R723SCS# contains a summary of the count and samples². This field can be used to apportion the service provided by the server to the various transaction service classes being served.

The following example illustrates the output from Rule WLM151:

RULE WLM151: SERVER SERVICE CLASS DELAYS

The CICUSRTX Service Class was served by the CICSGRN Service Class. The CICSGRN Service Class experienced the following delays during the measurement intervals when the CICUSRTX Service Class missed its performance goal (the delays are shown relative to the EXECUTING time of CICSGRN). CICSGRN also served other service classes. The "PCT SERVED" column reflects the percent of service provided by CICSGRN to CICUSRTX, relative to the total service provided by CICSGRN to all service classes served by CICSGRN.

MEASUREMENT INTERVAL	PCT CPU	PCT CPU	PCT CPU	PCT	PCT	PCT
	USING	DELAYED	CAPPING	PAGING WAIT	UNKNOWN WAIT	SERVED
13:07-13:12,21JUN1994	39.6	60.4	0.0	0.0	0.0	99.5
13:17-13:22,21JUN1994	42.3	57.7	0.0	0.0	0.0	99.7

Suggestion: There are no suggestions directly associated with this finding. CPExpert will continue analysis of the server service class(es), and other rules should be produced to provide suggestions.

²SMF field R723SCS# was improperly described in early versions of the SMF manual. The field has been modified to conform with the above description after CPExpert advised IBM of the error.